

PATENT
450110-03719

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

TITLE: IMAGE PROCESSING APPARATUS

INVENTORS: Stephen Mark KEATING, Jason PELLY, Ian
McLEAN

William S. Frommer
Registration No. 25,506
FROMMER LAWRENCE & HAUG LLP
745 Fifth Avenue
New York, New York 10151
Tel. (212) 588-0800

1006329-12001

5 Background of Invention

5 Background of Invention

Steganography is a technical field relating to the embedding of data into material such as video material, audio material and data material in such a way that the data is imperceptible in the material.

10 material and data material. A watermark may be imperceptible or perceptible in the material.

15 may identify the owner of the material.

20 Watermarks may be “fragile” in that they are easily damaged by processing which is useful to detect attempts to remove the mark or process the material.

25 watermark degrades the image and the mark is preferably not removable by the customer. Visible watermarks are also used to determine the provenance of the material into which they are embedded.

30 into a transform domain and adding the data to the image in the transform domain. For

the example of video images, the Discrete Wavelet Transform of these images into the transform domain, provides a form of the image in which the data can be embedded. The data to be embedded can be combined with the wavelet coefficients of one of a plurality of sub-bands which are formed in the transform domain. Generally, the data to be embedded is arranged to modulate a predetermined data sequence such as a Pseudo Random Bit Sequence (PRBS). For example, each bit of the data to be embedded is arranged to modulate a copy of the PRBS, and this copy is then added into one of the sub-bands of the image in the transform domain. The image is then converted back to the spatial domain.

Generally it is desirable in any steganographic scheme to reduce the effect of the embedded data on the material information in which the data is embedded, whilst as far as possible increasing the likelihood of correctly detecting the embedded data. For the example of video images, this represents reducing as far as possible a perceptible degradation of the image, whilst increasing the likelihood of correctly detecting the data.

Summary of Invention

According to the present invention there is provided an image processing apparatus operable to embed data into an image, the image comprising a frame of image data, the frame comprising first and second image fields, the apparatus comprising a combining processor operable to introduce the data to be embedded into data symbols of at least one of the first and second fields of the image frame.

It is known, for example, that video images can be generated by performing an inter-laced horizontal scan of an image. As a result, two image fields are produced for each image frame.

Some forms of processing which are applied to images can affect the image frames more than the image fields. As such, by embedding data into the fields of an image, an improvement in the robustness of the embedded data is provided. The data may be embedded in the first or the second or both of the fields of the image, with the effect that a likelihood of correctly recovering the data from the image may be increased. Furthermore, images are viewed as frames of data, so that combining the

data with the fields of the image can reduce a perceivable effect of the embedded data on the image.

Although the data to be embedded may be combined with the fields of the image in the spatial domain, in preferred embodiments the combining processor may be operable to represent the data to be embedded in a transform domain form, and, in combination with a transform processor, to combine the data to be embedded with the first and/or second fields of the image in one of a transform domain form, the transform processor generating a transform domain form of the first and/or second fields, the data being combined with the first and second fields by the combining processor in the transform domain, or a spatial domain form of the first and/or second fields, the transform processor generating a spatial domain representation of the transform domain form of the data to be embedded, the data being combined with the first and/or second fields of the image by the combining processor in the spatial domain.

Introducing the data into the first and/or second fields of the image in the transform domain provides an advantage of reducing any perceivable effect of the embedded data on the image. This is because combining the data with the transform domain symbols, effectively spreads the effect of the combined data over a greater number of symbols (pixels) in the spatial domain. The data can be combined with the image in either the spatial domain or the transform domain.

Advantageously, the transform domain may provide a plurality of sub-bands and the data to be embedded may be introduced into at least one of the sub-bands of the transform domain form of the first and/or second fields of the image. Furthermore, the combining processor may be operable in combination with the transform processor to combine the data to be embedded into one other sub-band of a transform domain form of the frame of the image data. The data may be combined with the other sub-band of the image frame in either the transform or the spatial domain.

As already explained, processing which may be performed on the image may affect the frames of the image more than the fields of an image. However, it is also possible that the fields of the image may be affected by noise and processing more than the frame of the image. Therefore by introducing the data to be embedded into a sub-band of the first and/or second image fields and introducing data into a different

sub-band of the image frame, some diversity is provided to the embedded data to the effect that a likelihood of correctly recovering the data may be increased.

In preferred embodiments the sub-band in which the data is introduced into one of the first and/or second fields, and the other sub-band in which the data is introduced into the image frame may include mutually exclusive spatial frequency components. Again, processing, noise and other disturbances to the image may affect one set of spatial frequency components more than another set of spatial frequency components. Therefore, by embedding the data into different sub-bands in the fields and frames of the image, a greater communication diversity is provided with the effect that the data may be more likely to be recovered correctly. Therefore, for example, the sub-band in which the data is introduced into the first and second fields may represent in the transform domain low spatial frequencies of the image in one direction and high spatial frequencies of the image in another direction, and the other sub-band in which the data is introduced in the image frame may represent in the transform domain high spatial frequencies of the image in the one direction and low spatial frequencies of the image in the another direction.

Although it will be appreciated that any appropriate transform can be used to embed the data into the image in the transform domain, in preferred embodiments, the transform is the discrete wavelet transform or the like. For the example of the discrete wavelet transform, the data symbols may be introduced into the first and/or second fields in one of the low vertical, high horizontal frequencies sub-band or the high vertical, low horizontal frequencies sub-band, and the other sub-band into which the data is introduced into the frame is the other of the low vertical, high horizontal frequencies and high vertical, low horizontal frequencies sub-bands.

Various further aspects and features of the present invention are defined in the appended claims.

Brief Description of the Drawings

Figure 1 is a schematic block diagram of a watermarking system;

Figure 2 is a schematic block diagram of a watermark embedder appearing in

Figure 1;

Figure 3 is a schematic representation illustrating the form of the wavelet domain transform;

Figure 4 is a schematic representation of an iterative filtering process which can be used to form the wavelet transform;

5 Figure 5 is a schematic representation illustrating the operation of the watermark embedder shown in Figure 2, at each of a plurality of stages A to F;

Figures 6a and 6b are graphical representations of an example frequency response of horizontal spatial frequencies of the image illustrating the effect of sub-sampling;

10 Figure 7 is a schematic illustration of the effect of sub-sampling on the frequency components of a wavelet sub-band;

Figure 8 is a schematic block diagram of the watermark decoder shown in Figure 1; and

15 Figure 9A and 9B are schematic block diagrams of the structure of an extended and a basic UMID respectively.

Description of Preferred Embodiment

An example embodiment of the present invention will be described with reference to a watermarking system in which data is embedded into a video image. Any type of data can be embedded into the image. However, advantageously the data
20 embedded into the image may be meta data which describes the image or identifies some attributes of the content of the image itself. An example of meta data is the Universal Material Identifier (UMID). A proposed structure for the UMID is disclosed in SMPTE Journal March 2000. A more detailed explanation of the structure of the UMID will be described later.

Watermarking System

Figure 1 illustrates a watermarking system, generally 10, for embedding a watermark into a video image 115, and recovering and removing a watermark from the video image 115. The watermarking system 10 in Figure 1 comprises an image
25 processor 100 for embedding the watermark into the video image, and a decoding image processor 102 for detecting and recovering the watermark, and for removing or
30 'washing' the watermark from the video image.

The image processor 100 for embedding the watermark into the video image comprises a strength adapter 180, and a watermark embedder 120. The watermark embedder 120 is arranged to embed the watermark into the video image 115, produced from the source 110, to form a watermarked image 125. The watermark to be
5 embedded into the video image is formed from data 175 representing a UMID. Generally, the UMID identifies the content of the video image, although it will be appreciated that other types of meta data which identify the content or other attributes of the image can be used to form the watermark. In preferred embodiments the watermark embedder 120 embeds the UMID into the video image 115 in accordance
10 with a particular application strength 185 from the strength adapter 180. The strength adapter 180 determines the magnitude of the watermark in relation to the video image 115, the application strength being determined such that the watermark may be recovered whilst minimising any effects which may be perceivable to a viewer of the watermarked image 125. After embedding the watermark, the image may be
15 transmitted, stored or further processed in some way, such as for example, compression encoding the image. This subsequent processing and transmitting is represented generally in Figure 1 as line 122.

In Figure 1 the decoding image processor 102 for detecting and removing the watermark is shown as comprising a watermark decoder 140, a data store 150 and a
20 watermark washer 130 which removes the watermark from the watermarked image 125.

The watermark decoder 140 detects the watermark from the watermarked video image and in the present example embodiment, generates a restored UMID 145 from the watermarked image 125. The watermark washer 130 generates a restored image
25 135, by removing as far as possible the watermark from the watermarked image 125. In some embodiments, the watermark washer 130 is operable to remove the watermark from the image substantially without leaving a trace. The restored image 125 may then be stored in a store 150, transmitted or routed for further processing.

The Watermark Embedder

The watermark embedder will now be described in more detail with reference to Figure 2, where parts also appearing in Figure 1 have the same numerical references.

- 5 In Figure 2 the watermark embedder 120 comprises generally a modulation processor 122 and a combining processor 124. Modulation processor 122 comprises a Pseudo Random Sequence Generator 220 which generates a pseudo random bit sequence which is fed to a first input of modulator 230. The UMID to be embedded into the image is received within the modulation processor 122 at an error correction encoder 200.

- 10 The error correction encoder operates to encode the UMID in accordance with an error correction encoding scheme. It will be appreciated that various error correction coding schemes could be used to encode the UMID. For the example embodiment the error correction encoder 200 utilises a Bose-Chaudhuri-Hocquenghem (BCH) systematic code providing 511 bit code words comprising 248 source bits of the UMID and 263 bits of redundant parity bits. Error correction encoding provides a facility for correcting errors which may be introduced into the encoded UMID, thereby improving the likelihood of correctly recovering the embedded UMID. However it will be appreciated that in other embodiments, the data may be combined with the image without being encoded using an error correction code.

- 20 The encoded UMID is fed to a second input of the modulator 230. In preferred embodiments the modulator 230 is operable to modulate each copy of a PRBS, generated by the pseudo-random sequence generator 220, with each bit of the error correction encoded UMID. The data is modulated by representing the values of each bit of the PRBS in bipolar form ('1' as +1, and '0' as -1) and then reversing the polarity of each bit of the PRBS, if the corresponding bit of the encoded UMID is a '0' and not reversing the polarity if the corresponding bit is a '1'. The modulated PRBS is then fed to the combining processor 124. The combining processor 124 will be described in more detail shortly.

- 30 Combining the encoded UMID with a PRBS has an effect of generating a spread spectrum signal representing the watermark data. In other embodiments the

UMID or error correction encoded UMID may be combined with the image without being spread spectrum encoded by modulating the PRBS with each bit of the data. However, in preferred embodiments the UMID or encoded UMID is spread spectrum encoded because this allows a reduction to be made in the strength of the data to be embedded in the image. By correlating the data in the transform domain image to which the modulated PRBS has been added, a correlation output signal is produced with a so called correlation coding gain which allows the modulated data bit to be detected and determined. As such, the strength of the data added to the image can be reduced, thereby reducing any perceivable effect on the spatial domain image. The use of a spread spectrum signal also provides an inherent improvement in robustness of the image because the data is spread across a larger number of transform domain data symbols.

In preferred embodiments, the combining processor 124 combines the PRBS modulated data with the image in different sub-bands of a wavelet transform of the image for both the fields of the image and the frame of the image. The encoded UMID is embedded in first and second fields of the image, and a copy of the encoded UMID is embedded in the frame of the image. However in alternative embodiments the encoded UMID may be divided into first and second parts, the first and second parts being embedded into the fields and frame of the image respectively. Generally, therefore in the following description the data embedded into the fields of the image will be referred to as a first data set, and the data embedded into the frame is referred to as a second data set.

The combining processor can be considered to be made up from a field embedding branch 261 comprising the processors in the bottom half of the combining processor 124 and a frame embedding branch comprising processors in the top half of the combining processor 124. The modulated PRBS data received from the combining processor 122 is combined with the image in the example embodiment of the present invention in the transform domain. Although any transform could be used, in the example embodiment the transform is the wavelet transform. In order to explain the embodiment of the present invention a brief description of wavelet transforms will be provided in the following paragraphs with reference to Figures 3 and 4.

The Wavelet Transform

Wavelets are well known and are described in for example "A Really Friendly
5 Guide to Wavelets" by C Valens, 1999 (c.valens@mindless.com). Valens shows that the discrete wavelet transform can be implemented as an iterated filter bank, as used in sub-band coding, with scaling of the image by a factor of 2 at each iteration.

Thus, referring to Figure 3 a spatial domain image is applied to a set of high
pass HP and low pass LP filters. At level 1, the first stage of filtering, the image is
10 filtered horizontally and vertically and, in each direction, scaled down by a factor of 2. In level 2, the low pass image from level 1 is filtered and scaled in the same way as in level 1. The filtering and scaling may be repeated in subsequent levels 3 onwards.

The result is shown schematically in Figure 4. Figure 4 is a representation
normal in the art. The horizontal axis H indicates increasing frequency. At level one
15 the image is spatially filtered into four bands; the lower horizontal and vertical band, I_{H1} , I_{V1} ; the upper horizontal band hH_1 , I_{V1} ; the upper vertical band I_{H1} , hV_1 ; and the upper horizontal and vertical band, hH_1 , hV_1 . At level 2, the lower horizontal and vertical band, I_{H1} , I_{V1} is filtered and scaled into the lower horizontal and vertical band, I_{H2} , I_{V2} ; the upper horizontal band hH_2 , I_{V2} ; the upper vertical band I_{H2} , hV_2 ; and the
20 upper horizontal and vertical band, hH_2 , hV_2 . At level 3 (not shown in Figure 4), the lower horizontal and vertical band, I_{H2} , I_{V2} is further filtered and scaled.

Combining Processor

The operation of the combining processor 124 shown in Figure 2 will now be
explained in more detail with reference to Figure 5. Figure 5 provides an illustrative
25 representation of the effect of processes performed on the image by the combining processor 124. These are represented as separate stages A to F, which represent the order and effect of processes performed by the combining processor 124.

Referring to Figure 2, the image 115 is received from the source 110 at a field
processor 181. For the example embodiment it is assumed that the image is formed
30 from an interlaced scan of the image. Generally therefore the image will be comprised of two fields. The first and second fields FLD0, FLD1 are illustrated in Figure 5 at

stage A. The field processor 181 is arranged to separate the first and second fields FLD0, FLD1 as shown in Figure 5 and to feed the first and second fields to a wavelet transformer 210 via parallel conductors 211 as shown in Figure 2. The wavelet transformer 210 generates a first order wavelet transform of the first and second fields FLD0, FLD1. Correspondingly the wavelet transform for each field is shown in Figure 5, at stage B, which provides a corresponding representation of the first order wavelet transform of the first and second fields WT_FLD0, WT_FLD1 using a corresponding notation to that used in Figure 3.

Returning to Figure 2, the two wavelet transformed fields WT_FLD0, WT_FLD1 are fed to a first combiner 242 via parallel conductors 241. The first combiner 242 also receives at a second input the first set of modulated PRBS data corresponding to the first copy of the encoded UMID, and at a third input the application strength α with which the PRBS data is to be combined with the image frame. The combiner 242 then combines the modulated PRBS received from the modulation processor 122 from the modulator 230 with a selected sub-band of the wavelet transform of the first and second fields WT_FLD0, WT_FLD1. For the example embodiment illustrated in Figures 2 and 5, the data added to the fields is introduced into the low horizontal, high vertical frequencies sub-band IH_1hV_1 as represented by the hatched region in the wavelet transform of the first and second fields WT_FLD0, WT_FLD1. The first data set (first copy of encoded UMID) is further divided into first and second subsets of data, the first and second subsets being combined respectively with the first and second wavelet transformed fields WT_FLD0, WT_FLD1. The combiner 242 combines the modulated PRBS with the wavelet coefficients of the wavelet transformed fields WT_FLD0, WT_FLD1, by adding for each bit of the modulated PRBS a factor α scaled by ± 1 in dependence upon the value of the bit of the first or second sets of data. This is represented by equation (1).

$$X'_i = X_i + \alpha_n W_i \quad (1)$$

Where X_i is the i -th wavelet coefficient, α_n is the strength for the n -th PRBS and W_n is the n -th bit of the watermark to be embedded in bipolar form, which for the example embodiment is the bit of the error correction encoded UMID.

Once the modulated PRBS's representing the first copy of the encoded UMID have been combined with the sub-band of each of the wavelet transformed fields WT_FLD0, WT_FLD1, the wavelet transform fields WT_FLD0, WT_FLD1 are fed to an inverse wavelet transformer 250 via parallel conductors 251. The inverse wavelet transform produced two fields of the image FLD0a and FLD0a having the first set of data combined with the image fields. The first and second fields embedded with the first data set are represented at stage C in Figure 5.

Although in the example embodiment the first set of data is divided into sub-sets, each being combined with one of the first or the second fields of the image, it will be appreciated that in other arrangements the first data set could be combined with either the first or the second field in accordance with a particular data embedding requirement. Similarly, if there is sufficient embedding capacity available, the same data by be combined with each of the first and second wavelet transformed fields, thereby providing redundancy which can be used to improve the likelihood of recovering the data.

The field embedded image is then fed to a frame processor 182 forming the first processor in the upper frame combining branch 263 of the combining processor 124. Therefore, as shown in Figure 5, at stage D, the spatial domain image having the first set of data embedded in the first and second fields is fed to the frame processor 182 which forms the fields into a frame.

Returning to Figure 2, the frame processor 182 feeds the image frame I_FRAME to a wavelet transformer 212 which generates as before a first order wavelet transform of the image frame. The wavelet transform of the image frame is then fed to a second combiner 240. The second combiner 240 also receives at a second input the second set of modulated PRBS data corresponding to the second copy of the encoded UMID, and at a third input the application strength α with which the PRBS data is to be combined with the image frame. The second set of data (second copy of encoded UMID) is then combined with the wavelet transform frame in a different wavelet sub-band than that to which the first set of data was added to the fields of the image. In accordance with the example embodiment as shown in stage E in Figure 5, the data is combined with the wavelet transform of the image frame WT_FRM in the high horizontal, low vertical frequencies sub-band $hH_1|V_1$. The data is combined, as

before, with the wavelet coefficients of the sub-band, as explained above with reference to equation (1).

Having added the second set of data into the high horizontal, low vertical frequencies sub-band hH_1lV_1 of the wavelet transform of the image frame WT_FRM , the second combiner 240 feeds the wavelet transformed image frame WT_FRM to an inverse wavelet transformer 252 which generates a spatial domain representation of the image frame to form the watermarked image presented at an output 125. Thus in Figure 5 the watermarked image I_FRAME' generated at the output of the inverse wavelet transformer 252 is formed at state F.

By embedding the data in only the two sub-bands hH_1lV_1 , lH_1hV_1 , the likelihood of detecting the embedded data is improved whilst the effects that the embedded data will have on the resulting image are reduced. This is because the wavelet coefficients of the high horizontal, high vertical frequencies sub-bands hH_1hV_1 are more likely to be disturbed, by for example compression encoding.

Compression encoding processes such as JPEG (Joint Photographic Experts Group) operate to compress images by modifying the high frequency components of the image. Therefore, writing the data into this sub-band hH_1hV_1 would reduce the likelihood of being able to recover the embedded data. Conversely, data is also not written into the low vertical, low horizontal frequencies sub-band lH_1lV_1 . This is

because the human eye is more sensitive to the low frequency components of the image. Therefore, writing the data in the low vertical, low horizontal frequencies sub-band would have a more disturbing effect on the image. As a compromise the data is added into the high horizontal, low vertical frequencies sub-band hH_1lV_1 and the low horizontal, high vertical frequencies sub-bands lH_1hV_1 . However in order to increase

the likelihood of correctly detecting the embedded data and correspondingly reducing the perceived effects of the embedded data on the image, the data is written into the image into these sub-bands in the same direction as the low frequency components of the image. That is to say, the scan direction SV, SH is in the same direction as the spatial direction of the low frequencies of the wavelet sub-band. As mentioned this

provides a particular advantage by increasing the likelihood of correctly detecting the embedded data because the data is being written in a direction which corresponds to

the lower energy components of the image. Correspondingly, the strength of the embedded data can be reduced such that the effect on the image in the spatial domain can be reduced.

In preferred embodiments, the field processor 181 operates to sub-sample the data representing the first and second fields to the effect of shifting the sampling frequency and therefore the Nyquist frequency of the sampled data. The combining processor 240 therefore combines the modulated PRBSs of the first data set with the wavelet coefficients formed from the sub-sampled version of the first and second fields in the wavelet transform domain. Therefore, when half of the horizontal components of the image in one field are sub-sampled, the frequency response of the components of the image shown in Figure 6a is converted as shown in Figure 6b providing to some extent an improved resistance to aliasing errors when the data is combined with the sub-band. To this end, the field processor 181 performs the sub-sampling by picking every other line of the image frame to form the first of the two fields FLD0, FLD1, the remaining lines forming the second of the two fields. The first combiner 240 combines the first data set with the low horizontal, high vertical frequencies IH_1hV_1 sub-band in the first and second fields in this form. As a result, the effect on combining the data in the sub-sampled domain is represented in Figure 7. As indicated by arrow 302 the information in the low horizontal, high vertical frequencies sub-band IH_1hV_1 has shifted more towards the centre of the wavelet domain representation of the image. This has an effect of providing some increased resistance to compression encoding because, as already explained, the high frequency components of the image are more likely to be corrupted.

The example embodiment of the present invention provides an improved robustness with respect to the likelihood of correctly recovering the embedded data and also reducing the likelihood of the embedded data having an effect on the image. This is because processes performed on the image and other noise can affect the field components and the frame components of the image differently. It is for this reason that, in preferred embodiments, the data to be embedded is copied, and the two copies, which form the first and second data sets, are embedded into the fields of the image and the frame of the image data, respectively. It is therefore less likely that both versions of the encoded UMID will be affected in the same way in both the frame of

image data and the fields of the image. Some resistance to the disturbing effects of processing and noise is therefore provided as a result of an embedding diversity provided by combining copies of the data into the fields and the frame separately. Furthermore, as illustrated in the example embodiment, the data is embedded into different sub-bands of the wavelet transform. As such some processing, for example, may affect the low horizontal components of the image frame more than the high horizontal components. Therefore, the data embedded into the high horizontal sub-band increases the likelihood of recovering the data embedded from this sub-band correctly. Correspondingly, combining the data with the low horizontal, high vertical frequencies sub-band of the field components also provides a further resistance to the disturbing effects of processing and where the effects are more likely to be present in the high horizontal frequency components. This is particularly so where the fields have been generated in the horizontal domain. For the present example, the first and second fields both represent half the horizontal information in the image. Therefore, embedding the data into the low horizontal, high vertical frequencies sub-band again provides some reduction in the likelihood that the data would be corrupted to the extent that it cannot be correctly recovered at a decoder.

The form of the watermark embedder and the order in which the first and second data items in the image fields and image frame respectively is not limiting and is provided to facilitate understanding. It will be appreciated that there are other ways of embedding the data. In alternative arrangements, the first and second data sets comprising different content may be embedded in the image frame and fields in parallel. Furthermore, it may not be necessary to generate an inverse wavelet transform between embedding the first and second data sets.

Although in this example embodiment, the data to be embedded is combined with the image in the transform domain, in alternative embodiments the data could be represented in the transform domain, inverse transformed into the spatial domain, and added to the image in the spatial domain.

Decoder

The operation of the watermark decoder 140 will now be explained in more detail, with reference to Figure 8, where parts also appearing in Figure 1, bear identical reference numerals.

The decoder 140 for detecting and recovering the data from the watermarked image is shown in more detail in Figure 8. The watermark decoder 140 is generally divided into two branches, the upper branch 410 is arranged to detect data embedded into the image frame whereas the lower branch 412 is arranged to detect data embedded into the fields of the image.

In preferred embodiments the upper and lower branches 410, 412 operate in parallel to generate, for each embedded bit of the encoded UMID a correlation signal from the data embedded in the fields and the frame of the image respectively. The correlation signal for each bit of the encoded UMID is then combined to form a composite signal from which the bit of the encoded UMID is recovered. The operation of the decoder 140, will now be described.

The watermarked image 125 is received at a frame store 305 forming the first processing element of the upper branch 410, as shown in Figure 8. The frame store 305 stores an image frame from which the data is to be recovered. A wavelet transformer 310 generates a wavelet transform of the image frame which is then fed to a correlator 330. The pseudo random sequence generator 320 is arranged to reproduce the PRBS which was used to embed the encoded UMID as explained with reference to the watermark embedder shown in Figures 1 and 2. The reproduced PRBS is fed to a second input of the correlator 330.

The correlator 330 generates a correlation output signal for each bit of the second set of data corresponding to the second copy of the encoded UMID, which has been combined with the high horizontal, low vertical frequencies sub-band $hH_1|V_1$ by correlating the coefficients of this sub-band with the reproduced PRBS. This correlation is expressed as equation (2), below, where X_n is the n -th wavelet coefficient and R_i is the i -th bit of the PRBS generated by the Pseudo Random Sequence Generator 320.

$$C_n = \sum_{i=1}^I X_{xy+i} R_i \quad (2)$$

Data is recovered from the fields of the image by separating the image into the first and second fields FLD0, FLD1. This is performed in the lower part 412 of the watermarked decoder 140 by a field processor 334 which divides the watermarked image frame into the first and second fields FLD0, FLD1. The first and second fields of the watermarked image are then fed to a second correlator 338. The second correlator 338 also receives the reproduced PRBS from the pseudo random sequence generator 320. Accordingly, it will be appreciated that each bit of the first set of data is correlated with the reproduced PRBS from the low horizontal, high vertical frequencies sub-band $1H_1hV_1$ from each of the first and second image fields to produce a correlation signal as explained with reference to data recovered from the image frame.

The correlation signals for each bit of the encoded UMID recovered respectively from the first and second correlators 330, 338 are fed to a combiner 332 and combined to form a composite signal. The data bits of the encoded UMID are then recovered from the relative sign of the composite signal, which gives an indication of the value of the bit of the embedded data in correspondence with the sign used to represent this bit in the watermark embedder.

The reproduced encoded UMID is then fed to an error correction decoder 350 which decodes the encoded UMID in accordance with the error correction code applied by the encoder 200 shown in Figure 2. At the output of the decoder 350 a recovered version of the UMID is produced.

It will be appreciated that for both the example implementations of the watermark embedder shown in Figure 2 and the watermarked decoder shown in Figure 8, the embedding of the data in the fields and the frames of the image could be done in a different order or indeed could have been done in parallel. The invention is therefore not limited to any particular arrangement in which the data is embedded in the fields or the frames of the image or indeed whether the embedding is performed in the transform domain or the spatial domain. In particular, for example, if the data was added in the spatial domain, then the first and second parts of the encoded UMID to be

embedded in the image could be performed in parallel for the image frame and the image fields respectively.

5

The Universal Material Identifier (UMID)

A brief explanation will now be given of the structure of the UMID, with reference to Figure 9A and 9B. The UMID is described in SMPTE Journal March 2000. Referring to Figures 9A an extended UMID is shown to comprise a first set of 10 32 bytes of a basic UMID, shown in Figure 9B and a second set of 32 bytes referred to as signature metadata. Thus the first set of 32 bytes of the extended UMID is the basic UMID. The components are:

- A 12-byte Universal Label to identify this as a SMPTE UMID. It defines the type of material which the UMID identifies and also defines the methods by which the 15 globally unique Material and locally unique Instance numbers are created.

- A 1-byte length value to define the length of the remaining part of the UMID.

- A 3-byte Instance number which is used to distinguish between different 'instances' of material with the same Material number.

- A 16-byte Material number which is used to identify each clip. Each Material 20 number is the same for related instances of the same material.

The second set of 32 bytes of the signature metadata as a set of packed metadata items used to create an extended UMID. The extended UMID comprises the basic UMID followed immediately by signature metadata which comprises:

- An 8-byte time/date code identifying the time and date of the Content Unit 25 creation.

- A 12-byte value which defines the spatial co-ordinates at the time of Content Unit creation.

- 3 groups of 4-byte codes which register the country, organisation and user codes.

30 More explanation of the UMID structure is provided in co-pending UK patent application number 0008432.7.

Various modifications may be made to the embodiments herein before described without departing from the scope of the present invention.